

Federated Community Cloud Standards

Current Standards Efforts and Status of Interoperability Work With Other SDOs

Alan Sill, Ph.D

Senior Scientist, High Performance Computing Center
Adjunct Professor of Physics, Texas Tech University
OGF Vice President of Standards

US NIST Federated Community Cloud Subgroup

March 30, 2012

OGF Standards Background



- OGF has extensive set of applicable standards related to federated community grid and cloud computing:
 - Federated Identity Management (FedSec-CG)
 - Managing the Trust Eco-System (CA operations, AuthN/AuthZ)
 - Virtual Organizations (VOMS)
 - Job Submission and Workflow Management (JSDL, BES)
 - Network Management (NSI, NML, NMC, NM)
 - Secure, fast multi--party data transfer (GridFTP, SRM)
 - Data Format Description (DFDL)
 - Service Agreements (WS-Agreement, WS-Agreement Negotiation)
 - Cloud Computing interfaces (OCCI)
 - Distributed resource management (DRMAA, SAGA, etc.)
 - Firewall Traversal (FiTP)
 - Others under development
- Working to gather this information to form an organized description of OGF work - an OGF “**Cloud Portfolio**”.

OGF and major science infrastructure projects - US and International



- OGF began in 2001 as an organization to promote the advancement of distributed computing worldwide.
- Grid Forum --> Global Grid Forum --> GGF + Enterprise Grid Alliance --> formation of OGF in 2005.
- Mandate is to take on all forms of distributed computing and to work to promote cooperation, information exchange, best practices use and standardization.
- OGF best known for a series of important grid standards that form the basis for major science and business-based distributed computing (BES, GridFTP, DRMAA, WS-Agreement, RNS, GLUE, etc.).
- Have been working on cloud standards (OCCI, etc.) for several years; several cooperative work agreements with other SDOs in place as appropriate also.

Recently Published OGF Standards



- OCCI core, infrastructure and RESTful HTTP rendering (GFD.183, 184 and 185 respectively).
- Data Management API within the GridRPC (GFD.186).
- OGSA-DMI Plain Web Service Rendering (GFD.187).
- WS-Iterator 1.0 (GFD.188)
- Relying Party Defined Namespace Constraints Policies in a Policy Bridge PKI Environment (GFD.189).
- Mapping between DFDL 1.0 Infoset and XML Data Model (GFD.190, supplements Data Format Description Language, GFD.174).
- Procedure for Registration of Subnamespace Identifiers in the URN:OGF Hierarchy (GFD.191).
- Web Services Agreement Specification (WS-Agreement) (GFD.192) and WS-Agreement Negotiation (GFD.193).
- Distributed Resource Management Application API Version 2 (DRMAA) (GFD.194, obsoletes 22, 130 & 133).

Recently Published OGF Standards



- OCCI core, infrastructure and RESTful HTTP rendering (GFD.183, 184 and 185 respectively).

**Takeaway message:
OGF continues to make
progress on its basic
core infrastructure
specifications**

- Distributed Resource Management Application API Version 2 (DRMAA) (GFD.194, obsoletes 22, 130 & 133).

OGF Cooperative Agreements In Place as of March 2012



OCCL and DMTF:

- OGF published the the OCCL Core, Infrastructure and HTTP Rendering specifications as GFD.183, 184 and 185 respectively, and is working on a JSON rendering. We created a joint work register with DMTF and continue to follow their progress towards publishing the related CIMI specification.

OGF and ISO:

- OGF has been accepted as a Category A liaison with ISO JTC1 SC38 on Cloud Computing.

OCCL and CDML:

- OGF has cooperative agreement w/SNIA and has held 4 jointly hosted Cloud Standards Plug-Fests so far; series continues.

Other Cooperative Agreements As of Now or On The Horizon:



OGF and TM Forum:

- Final stages of preparing cooperative agreement on Service Agreements, including SLAs, between OGF and TM Forum. TM Forum has an agreement in place with DMTF, and is contributing to NIST discussions on SLAs.

OGF and CSA:

- Cooperative agreement between OGF and CSA in place.

OGF and IEEE:

- OGF co-sponsored IEEE Cloudcom 2011 (Athens) and is involved in planning for Cloudcom 2012 (Taipei).

OGF and SIENA, NIST, GICTF, etc.:

- We to contribute to roadmap processes and work actively with these to contribute to ongoing roadmapping efforts.

XSEDE: The Next Generation of US Supercomputing Infrastructure

The Role of Standards for Risk Reduction and Inter-operation in XSEDE

Andrew Grimshaw

XSEDE

Extreme Science and Engineering
Discovery Environment

XSEDE Services Layer:

LSN-MAGIC Meeting
February 22, 2012

Simple services combined in many ways

- Resource Namespace Service 1.1
- OGSA Basic Execution Service
- OGSA WSRF BP – metadata and notification
- OGSA-BytelIO
- GridFTP
- WS Trust Secure Token Services
- WSI BSP for transport of credentials
- ... *(more than we have room to cover here)*

Examples – (not
a complete list)

Basic message (AFS): XSEDE represents a phase change in the engagement of OGF standards with US cyberinfrastructure.

Why Open Standards?

LSN-MAGIC Meeting
February 22, 2012

- *Risk reduction*
- *Best-of-breed mix-and-match*
- *Allows innovation/competition at more interesting layers*
- *Facilitates interoperation with other infrastructures*

Takeaway message

- *The use of standards permits XSEDE to interoperate with other infrastructures, reduces risks including vendor lock-in, and allows us to focus on higher level capabilities and less on the mundane*

UNICORE in XSEDE: Towards a large-scale scientific Environment based on Open Standards

Evolution from TeraGrid

Starting in 2001, the National Science Foundation program TeraGrid has developed into one of the world's largest and most comprehensive Grid projects, offering resources and services to more than 10,000 scientists. Its successor, the Extreme Science and Engineering Discovery Environment (XSEDE, www.xsede.org), has started in July 2011 and is expected to excel the previous program in terms of service quality while lowering technological entry barriers at the same time. These and other goals are to be achieved in the project's five year grant period with an overall budget of \$121 million. Among the partnership of 17 institutions, the Jülich Supercomputing Centre (JSC) is the only organization located outside the USA.

Open Standards-based Architecture

Since many scientific communities operate internationally, one key element of XSEDE is the use of open standards in order to promote interoperability with other distributed computing infrastructures such as PRACE in Europe. Figure 1 shows the extended reference architecture providing mandatory XSEDE Enterprise Services at every major XSEDE site as well as optionally available Community Provided Services. For many years, the JSC and several other XSEDE partners have been active

in establishing the key standards that now define the interfaces of the XSEDE Enterprise Services. Within the Grid Interoperation Now (GIN) community group of the Open Grid Forum (OGF, www.ogf.org), such key standards as BES/JSDL for running remote computations have demonstrated their impact on scientific applications. Based on these standards, scientific workflows can be executed today across different infrastructures with no less than 8 different Grid middleware technologies.

Jülich's Role in XSEDE

The JSC not only contributes its extensive Grid know-how gained from European research projects and its experience in standard-based software engineering, but also a technology called Uniform Interface to Computing Resources (UNICORE, www.unicore.eu). Being developed by partners all over Europe, UNICORE is a Grid system that provides secure and seamless access mechanisms to a variety of different computer systems and platforms. It facilitates the remote execution of scientific applications as well as sharing software, resources and data. UNICORE is fully based on Web services and open standards in order to allow seamless interoperation with other standard compliant Grid systems such as Genesis II which is developed at the University of Virginia. Being complementary to the more

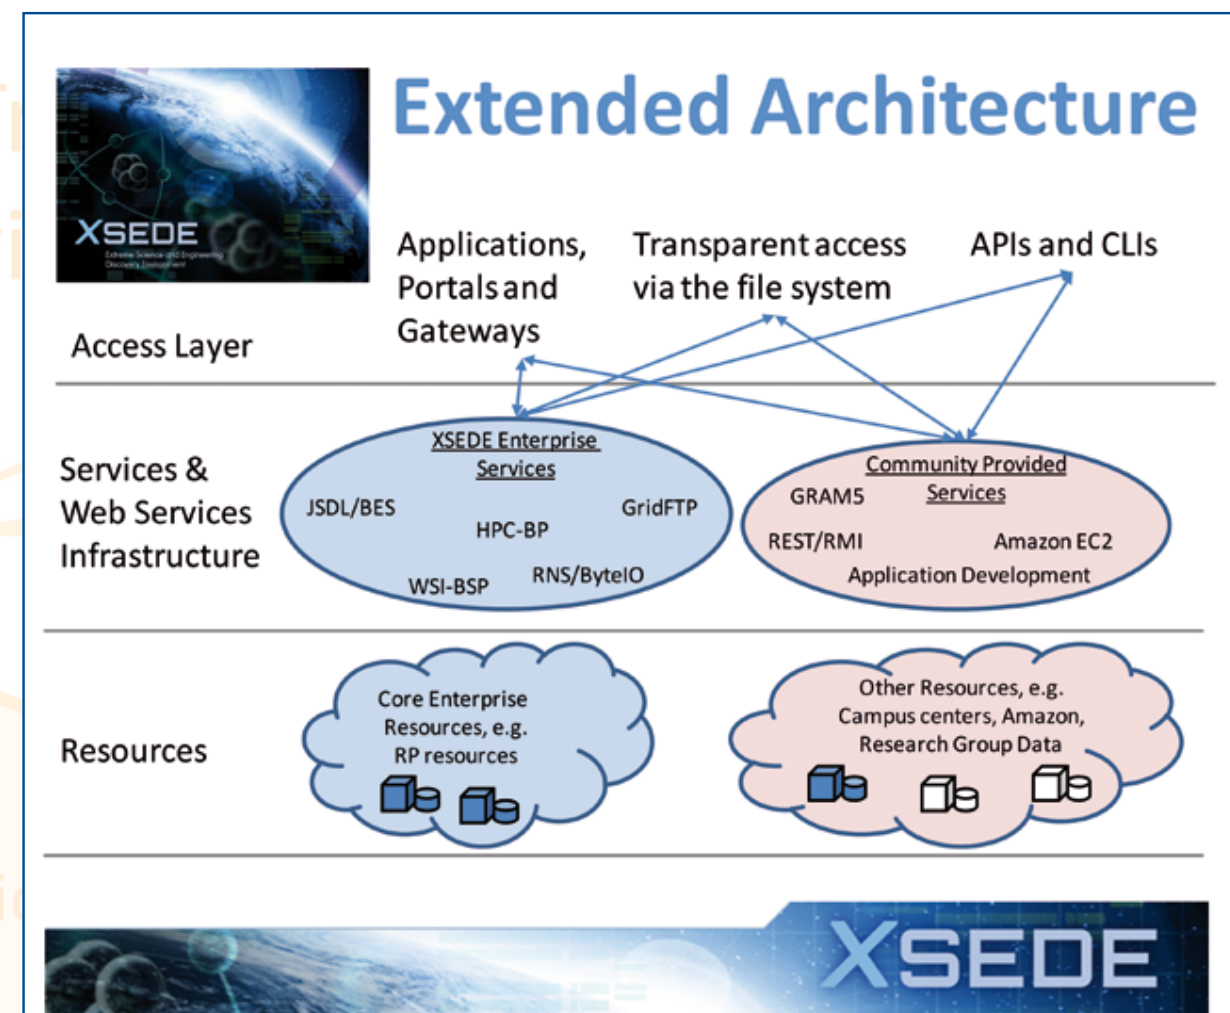


Figure 1: The current XSEDE architecture aims at providing XSEDE Enterprise Services at every major XSEDE site and optionally available Community Provided Services. The architecture will evolve over time according to end-users' needs.

• Morris Riedel
• Bastian Demuth

lightweight Genesis II services, UNICORE meets all the security requirements of modern High Performance Computing centres and provides extensive support for their highly specialized hardware as well as their varying batch systems.

Infrastructure Vision

The XSEDE architecture envisions deploying UNICORE as part of the XSEDE Enterprise Services at major US high performance centres whereas Genesis II will be used for integrating smaller computer systems such as desktop PCs in order to provide interoperability with campus Grids across the country. The

resulting infrastructure is expected to cover both high performance and high throughput computing, thus enabling innovative research and discovery requiring both types of parallel computations. Moreover, collaboration between American and European scientists will be easier than ever: UNICORE will also be deployed on the European Grid Infrastructure (EGI) and is already installed on many of the systems forming the infrastructure of the European supercomputing project PRACE.

Jülich
Supercomputing
Centre

UNICORE in XSEDE: Towards a large-scale scientific Environment based on Open Standards

inSiDE

inSiDE • Vol. 9 No.2 • Autumn 2011

Innovatives Supercomputing in Deutschland




Figure 1: The current XSEDE architecture aims at providing XSEDE Enterprise Services at every major XSEDE site and optionally available Community Provided Services. The architecture will evolve over time according to end-users' needs.

• Morris Riedel
• Bastian Demuth

lightweight Genesis II services, UNICORE meets all the security requirements of modern High Performance Computing centres and provides extensive support for their highly specialized hardware as well as their varying batch systems.

Infrastructure Vision

The XSEDE architecture envisions deploying UNICORE as part of the XSEDE Enterprise Services at major US high performance centres whereas Genesis II will be used for integrating smaller computer systems such as desktop PCs in order to provide interoperability with campus Grids across the country. The

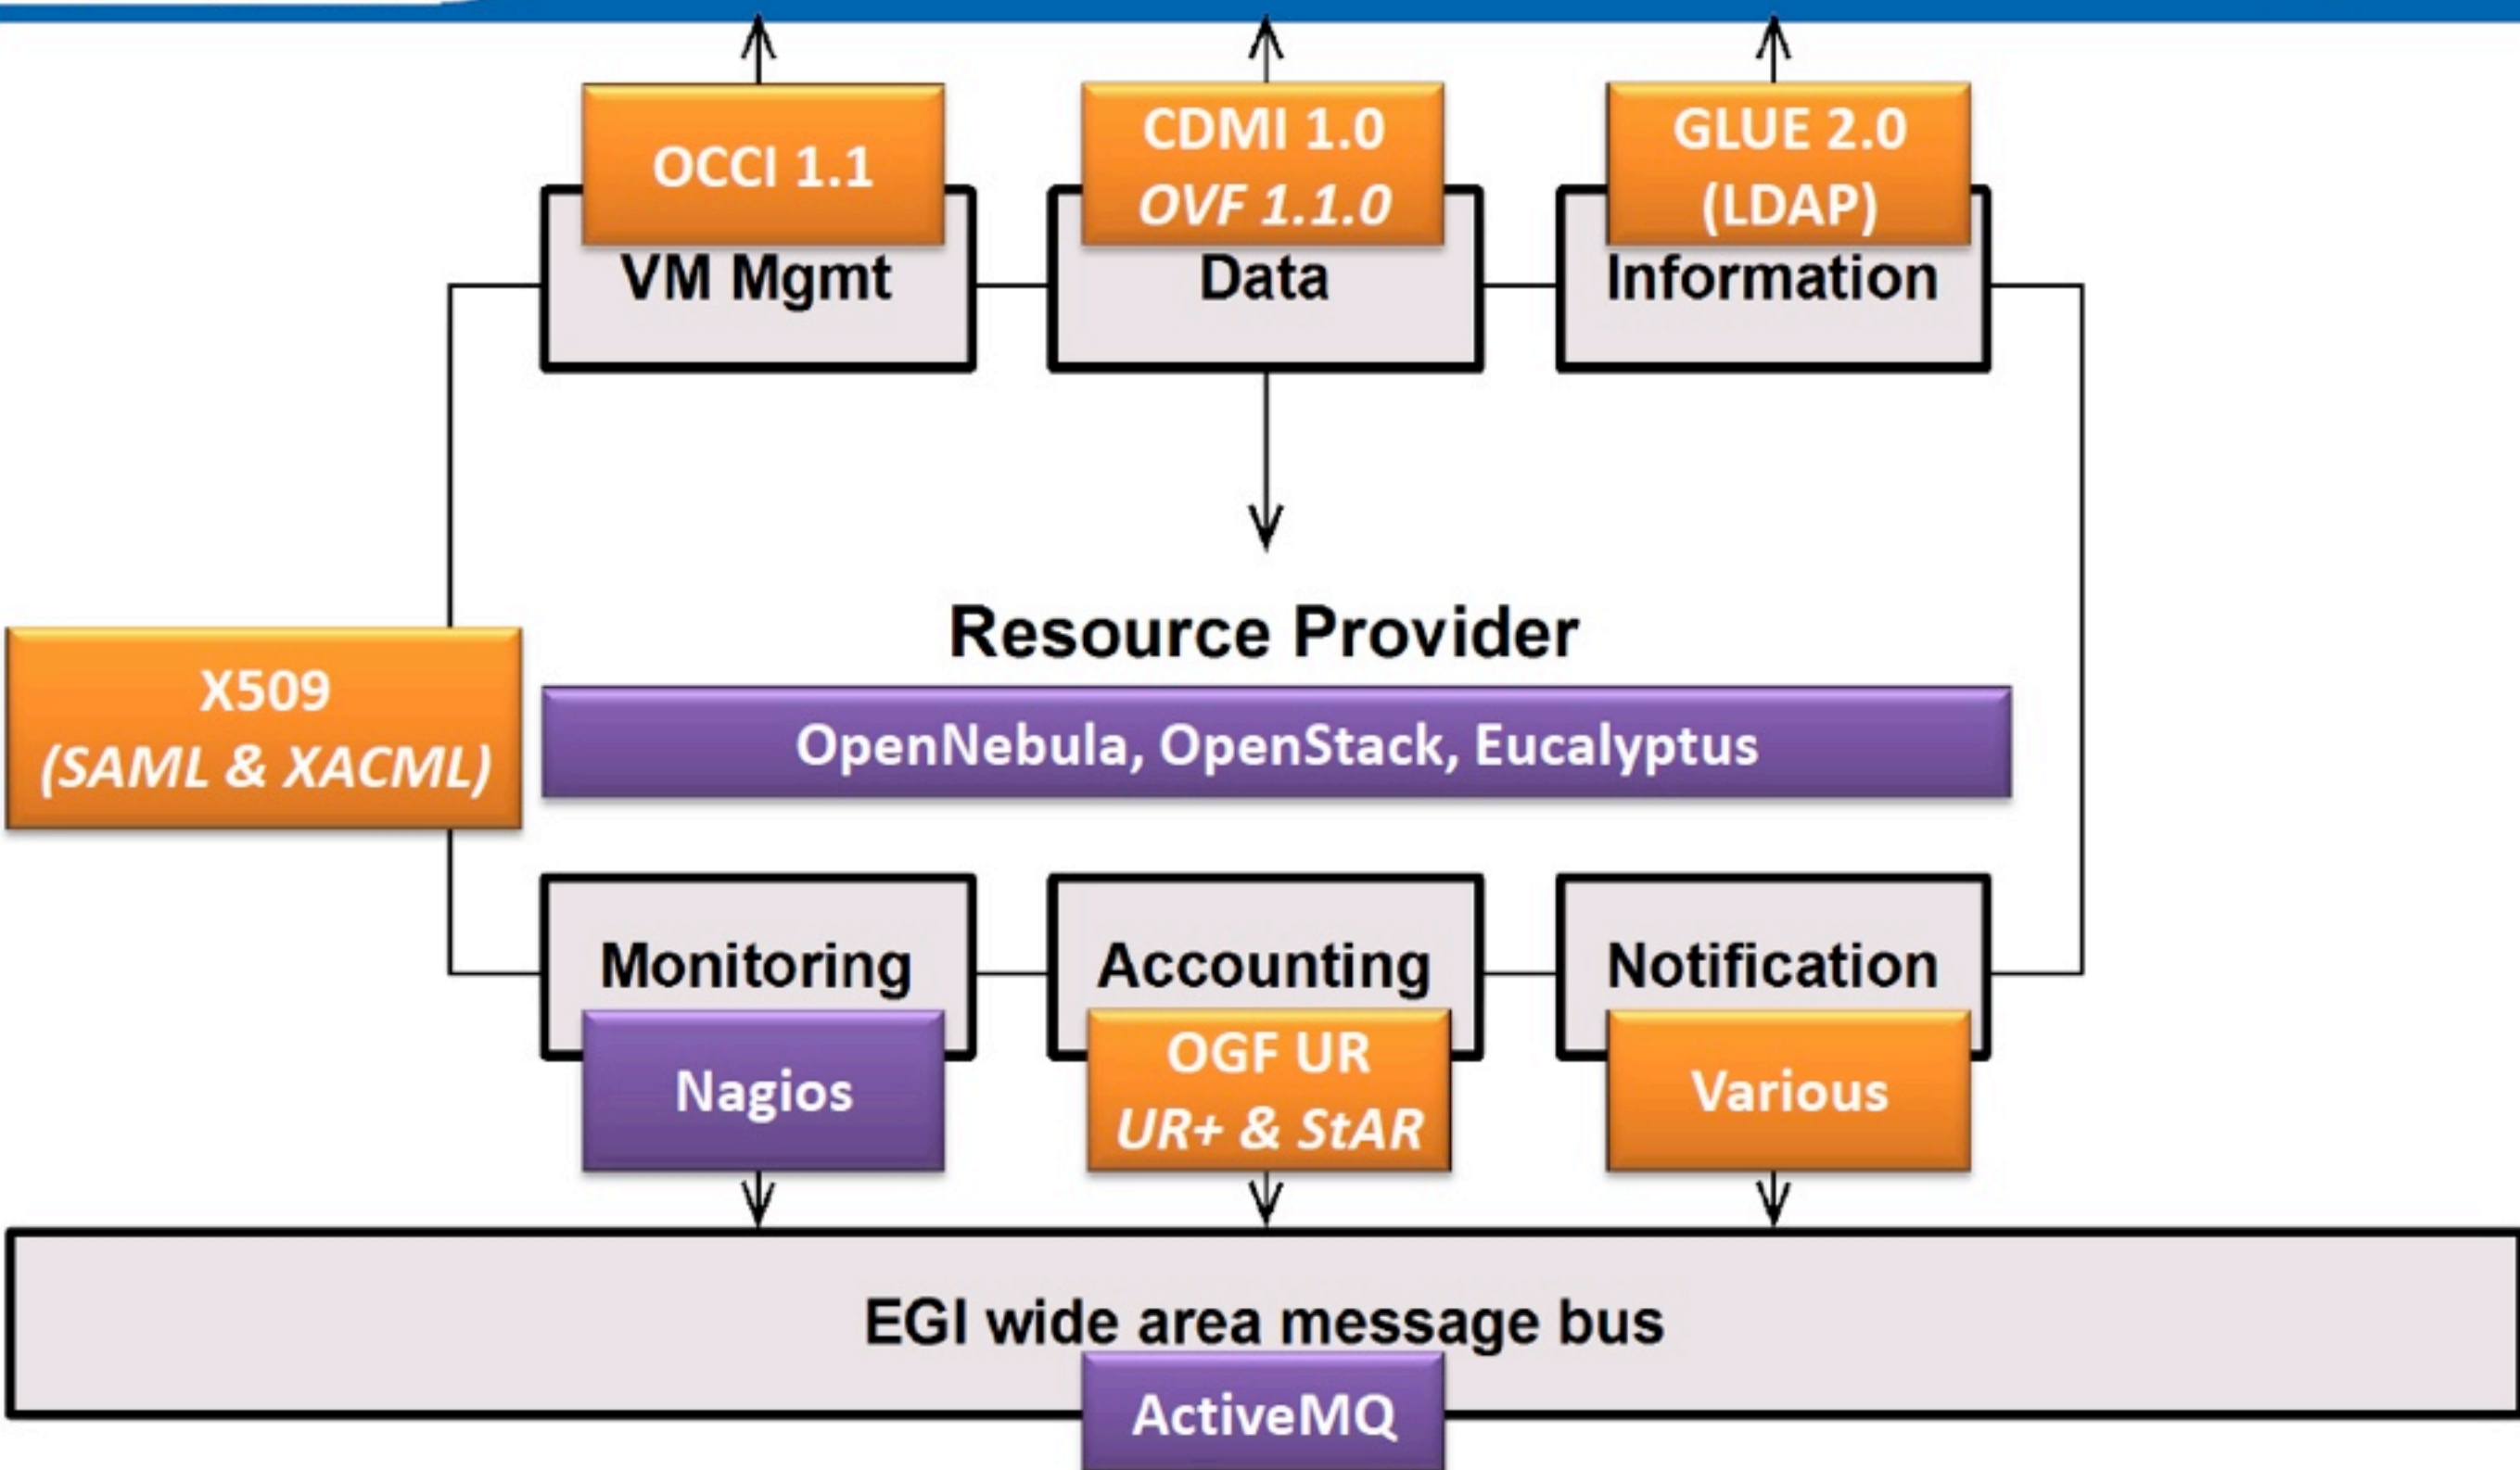
resulting infrastructure is expected to cover both high performance and high throughput computing, thus enabling innovative research and discovery requiring both types of parallel computations. Moreover, collaboration between American and European scientists will be easier than ever: UNICORE will also be deployed on the European Grid Infrastructure (EGI) and is already installed on many of the systems forming the infrastructure of the European supercomputing project PRACE.

Jülich
Supercomputing
Centre

European Grid Initiative Federated Clouds Task Force



- Objectives:
 - Write a blueprint document for resource providers that wish to securely federate and share virtualised environments as part of the EGI production infrastructure;
 - Deploy a test bed to evaluate integration of virtualised resources within the existing EGI production infrastructure for monitoring, accounting and information services;
 - Investigate and catalogue the requirements for community facing services based on or deployed through virtualised resources;
 - Provide feedback to relevant technology providers on their implementations and any changes needed for deployment into the production infrastructure;
 - Identify and work with user communities willing to be early adopters of the test bed infrastructure to help prioritise its future development;
 - Identify issues that need to be addressed by other areas of EGI (e.g. policy, operations, support & dissemination).
- <https://wiki.egi.eu/wiki/Fedcloud-tf:FederatedCloudsTaskForce>



European Grid Initiative Federated Clouds Testbed:



Cloud Resources Status

The Task Force is developing a resource monitor solution for the clouds federation based on Nagios. Meanwhile, here a table showing the current status of the cloud resources made available by the resource providers that have joined the Task Force. This table is updated weekly by the resource providers.

Providers	<div>■ = Available</div> <div>■ = Not available</div>	User registration	User access	VM availability	Elastic IPs	Object Storage	Persistent Storage
CESGA (IBERgrid)							
CESNET (NGI CZ)			OCCI v0.8 OCCI v1.1	VM suse (storage/16) with NET public (network/4)	No	Cumulus at carach3.ics.muni.cz:8888	GridFTP at carach4.ics.muni.cz:50000
CYFRONET (NGI PL)							
GWDG							
FZ Jülich		Mail to Björn Hagemeier	<ul style="list-style-type: none"> EC2 egi-cloud.zam.kfa-juelich.de:8773 S3 egi-cloud.zam.kfa-juelich.de:3333 		134.94.32.33 - 134.94.32.40		
IGI							
IN2P3 (NGI FR)							
KTH							
OerC (UK NGI)							
SARA (NGI NL)							

FutureGrid: (Potential augmentation of / replacement for SAJACC)



- Cloud Interoperability Testbed project hosted on the US FutureGrid project infrastructure
 - Cloud Infrastructure:
 - Nimbus
 - Eucalyptus
 - Openstack
 - Cloud Platforms:
 - Hadoop
 - Pegasus
 - Twister
 - Grid Platforms:
 - Unicore
 - Genesis II
 - HPC Tools:
 - Torque/Moab
 - MPI
 - ScaleMP
- CI-TB allows anyone interested in signing up to test standards, implementations for interoperability
- <https://portal.futuregrid.org/project/176>

FutureGrid: (Potential augmentation of / replacement for SAJACC)



Cloud Interoperability Testbed

Submitted by alansill on Tue, 29 Nov 2011, 13:16:12 GMT

Keywords:

open cloud computing interface OCCI

Project Contact

Provide information about who is the Project Lead and who will be interacting with FG staff in case of communications. For the later we provide the option to specify a project manager. In many cases the Lead and Manager will be the same.

Project Lead:

alansill

Project Manager:

alansill

Project Contact:

Alan.Sill@ttu.edu

Project Information

ProjectId:

176

Project Orientation:

Industry

Primary Discipline:

Computer Science (401)

Abstract:

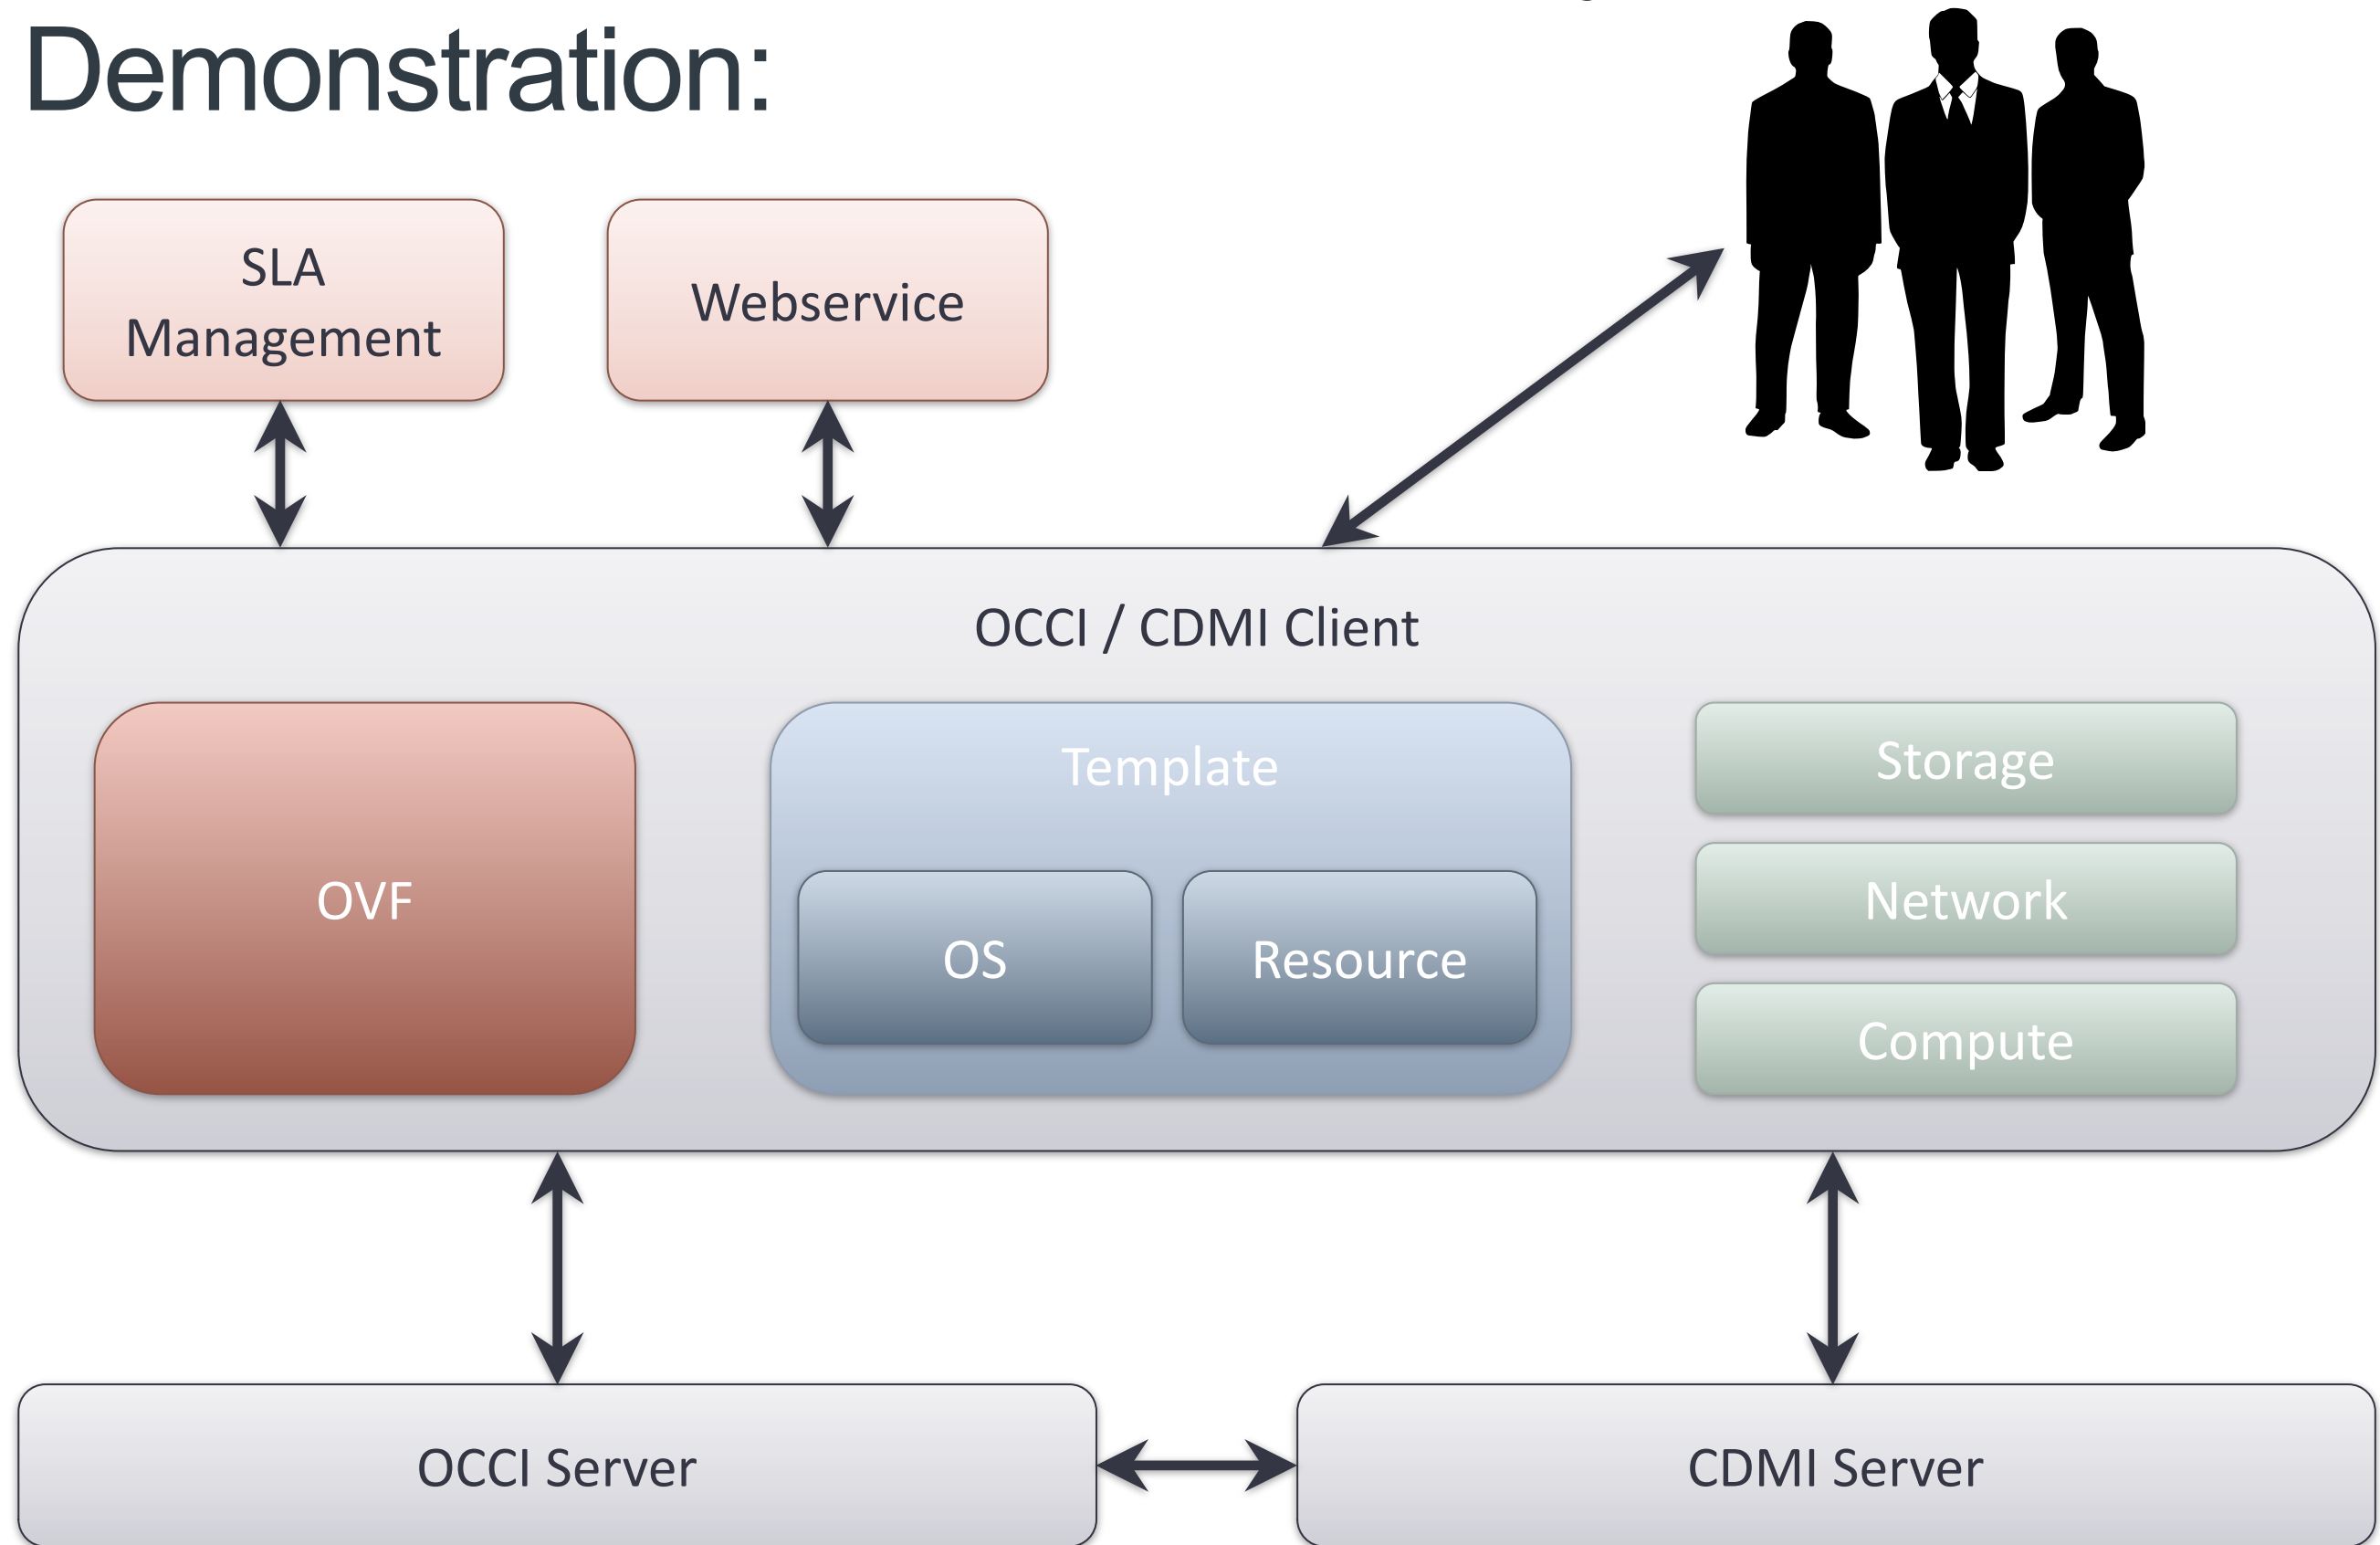
The Cloud Interoperability Testbed will serve as a mechanism to host interoperability tests for different machine control, data transfer, resource reporting and usage agreement standards and implementations of other new standards efforts. The primary purpose will be to give developers an opportunity to try out implementations of code that implements either server or client functionality for the use of multiple standards from different standards development organizations.

Other Cloud Interoperability Efforts



- SNIA “Cloud Plugfests” being used for testing combined interoperability efforts: April, July, Sept 2011, Spring 2012.
- *Several active OCCI and CDMI implementations!*
 - Progressing well into production ready solutions.
- Integration of OCCI and CDMI reportedly “works well.”
- OVF integration under active development & progressing.
- Next steps and new testing work planned:
 - Feed back experience from implementations, write implementation experience documents and user guides.
 - Develop combined OCCI / CDMI / OVF verification suite.
- Activities in IEEE p2301, p2302 just beginning (see below).
- OCCI currently awaiting publication of CIMI by DMTF for comparison and interoperability studies
- New Cloud Interoperability group launches at OGF 34!

OCCI/CDMI/OVF Interoperability Demonstration:



OCCI/CD Demonst



CLOUD PLUGFEST 2012 - YOUR INVITATION TO PARTICIPATE Registration Deadline is February 24, 2012

The Storage Networking Industry Association's Cloud Storage Initiative (SNIA CSI) and the Open Grid Forum (OGF) are conducting the first in a series of international cloud plugfests slated for 2012.

What: Cloud Plugfest sponsored by SNIA CSI/OGF

When: February 28 - March 1, 2012

Where: Dusseldorf, DE (Primary Location)

Who's invited?

All cloud implementers are invited to participate either in person or remotely. Participants need not be members of the SNIA, CSI or OGF and the **cost to attend is FREE!***

Why attend?

The cloud plugfest will offer a highly collaborative, vendor-neutral, environment for developers and vendors to perform interoperability testing of CDMI and OCCI implementations. The previous Cloud Plugfest, held in September 2011 in Santa Clara, California and remotely in Dusseldorf, Germany, attracted over 20 developers from nine companies.

By attending the cloud plugfest participants will:

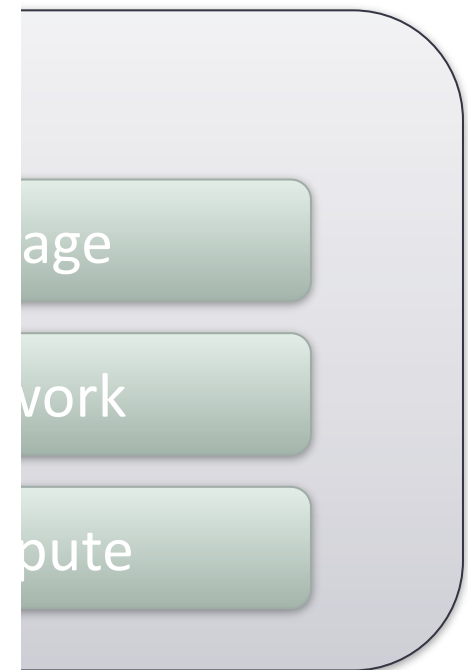

- Gain a greater understanding of what the needs are for establishing better interoperability with Cloud Computing and Cloud Storage standards, and the opportunity to refine those standards.
- Gain a better understanding of the requirements for interoperability integration between Cloud Storage (CDMI) and Cloud Computing (OCCI)
- Interact directly with cloud implementers and early adopters of CDMI and OCCI

How to register?

The cost to register is free. Visit the cloud plugfest page at www.snia.org/cloud/cloudplugfest for additional details.

* Plugfest registration is free. Attendees are responsible for all other costs (travel, meals, etc.) associated with participating in this event including any participation fees imposed by co-located venues or event sponsors.

For questions or comments about the SNIA Cloud Storage Initiative or this plugfest, please contact Tom Mancuso at csimanager@snia.org.



Updates On Recent and Updated OCCI Implementations



- New extensive commercial implementation of OCCI and other standards-based technologies from **CompatibleOne**:
<http://www.compatibleone.org>
- New **OpenNebula** implementation in progress - full OCCI v1.1 compliance! **OCCI Ruby client**, instructions and tutorials on developer site, including examples and testing tools available.
- Uptake of OCCI work by FP7 “Future Internet” **FI-Ware project**.
- **occi-py** Python framework developed by Ralf Nyren: RDBMs-friendly backend; draft JSON rendering implemented and demo instance available for testing.
- **R2AD**: New Android and JavaFX Cloud Client apps including OCCI and CDMI: design done, demos available and beta implementation in progress.
- **OpenStack** implementation code work nearly complete.
- New client and server libraries for several languages available.

More information

- OGF Standards:
 - <http://ogf.org/standards>
 - <http://ogf.org/documents>
- OCCI WG website: <http://www.occi-wg.org>
- DFDL Introduction: <http://www.ogf.org/dfdl/>
- WS-Agreement specification (obsoletes GFD.107):
<http://www.ogf.org/documents/GFD.192.pdf>
- WS-Agreement Negotiation specification:
<http://www.ogf.org/documents/GFD.193.pdf>
- DRMAA v2: <http://www.ogf.org/documents/GFD.194.pdf>
- Firewall Traversal specification (FiTP): (in progress)
- ... *Stay tuned! Others coming soon!*

E-mail:

standards@ogf.org

Full Copyright Text



Copyright (C) Open Grid Forum (2011). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works.

The limited permissions granted above are perpetual and will not be revoked by the OGF or its successors or assignees.